APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/531,397	03/21/2000	Joseph C. Ballantyne	3797.81466	6866

28319   7590   04/03/2003

BANNER & WITCOFF LTD.,
ATTORNEYS FOR MICROSOFT
1001 G STREET , N.W.
ELEVENTH STREET
WASHINGTON, DC   20001-4597

EXAMINER
ALI, SYED J

ART UNIT	PAPER NUMBER
2127	7

DATE MAILED: 04/03/2003

Please find below and/or attached an Office communication concerning this application or proceeding.

PTO-90C (Rev. 07-01)

	Application No.	Applicant(s)
Office Action Summary	09/531,397	BALLANTYNE, JOSEPH C.
	Examiner	Art Unit
	Syed J Ali	2127

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE <u>3</u> MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.
- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133).
- Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1)☒ Responsive to communication(s) filed on *March 21, 2000* .

2a)☐ This action is **FINAL**. 2b)☒ This action is non-final.

3)☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4)☒ Claim(s) *1-33* is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5)☐ Claim(s) _____ is/are allowed.

6)☒ Claim(s) *1-33* is/are rejected.

7)☐ Claim(s) _____ is/are objected to.

8)☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9)☐ The specification is objected to by the Examiner.

10)☐ The drawing(s) filed on _____ is/are: a)☐ accepted or b)☐ objected to by the Examiner.

Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).

11)☐ The proposed drawing correction filed on _____ is: a)☐ approved b)☐ disapproved by the Examiner.

If approved, corrected drawings are required in reply to this Office action.

12)☐ The oath or declaration is objected to by the Examiner.

Priority under 35 U.S.C. §§ 119 and 120

13)☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a)☐ All b)☐ Some * c)☐ None of:

1.☐ Certified copies of the priority documents have been received.

2.☐ Certified copies of the priority documents have been received in Application No. _____ .

3.☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

14)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. § 119(e) (to a provisional application).

a) ☐ The translation of the foreign language provisional application has been received.

15)☐ Acknowledgment is made of a claim for domestic priority under 35 U.S.C. §§ 120 and/or 121.

Attachment(s)

1)☒ Notice of References Cited (PTO-892)
2)☒ Notice of Draftsperson's Patent Drawing Review (PTO-948)
3)☒ Information Disclosure Statement(s) (PTO-1449) Paper No(s) <u>6</u> .
4)☐ Interview Summary (PTO-413) Paper No(s). _____ .
5)☐ Notice of Informal Patent Application (PTO-152)
6)☐ Other: .

DETAILED ACTION

Claim Rejections - 35 USC § 102

1.   The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the

basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in a patent granted on an application for patent by another filed in the United States before the invention thereof by the applicant for patent, or on an international application by another who has fulfilled the requirements of paragraphs (1), (2), and (4) of section 371(c) of this title before the invention thereof by the applicant for patent.

The changes made to 35 U.S.C. 102(e) by the American Inventors Protection Act of 1999

(AIPA) and the Intellectual Property and High Technology Technical Amendments Act of 2002

do not apply when the reference is a U.S. patent resulting directly or indirectly from an

international application filed before November 29, 2000. Therefore, the prior art date of the

reference is determined under 35 U.S.C. 102(e) prior to the amendment by the AIPA (pre-AIPA

35 U.S.C. 102(e)).

2.   Claims 1-3, 6-7, 9-10, 19-21, 24-25, and 33 are rejected under 35 U.S.C. 102(e) as being

anticipated by Reiffin (USPN 6,330,583) (hereinafter Reiffin).

As per claim 1, Reiffin teaches a method of scheduling CPU resources comprising the

steps of:

a.   using a counter to determine when to allocate CPU resources (col. 4 lines 1-14,

"When the end of the current timeslice is thereby determined by the counter the latter

then signals the interrupt controller which in turn activates the interrupt input of the

CPU");

b.   instructing an interrupt controller to allocate the CPU resources (col. 4 lines 1-14,

"When the end of the current timeslice is thereby determined by the counter the latter

then signals the interrupt controller which in turn activates the interrupt input of the

CPU"); and

c.   allocating the CPU resources (col. 4 lines 33-40, "The scheduler signals its task

selection to the CPU", wherein by selecting a task, the scheduler tells the CPU where to

allocate resources).

As per claim 2, Reiffin teaches the method of claim 1 wherein only a portion of the CPU

resources are allocated (col. 4 lines 41-59, "control of the CPU is rapidly switched back and forth

between a local task and a network task", wherein since multiple tasks are running concurrently,

the task being serviced as a result of the interrupt is only occupying a portion of the CPU

resources).

As per claim 3, Reiffin teaches the method of claim 1 wherein all of the CPU resources

are allocated (col. 6 lines 14-24, "Execution of either the resumed network task or the newly

fetched network task continues until either the task is completed or...the tasks is preempted by

the termination of the timeslice", wherein the network occupies the entirety of the CPU resources

until it is either completed or the timeslice expires).

As per claim 6, Reiffin teaches the method of claim 1 wherein the counter is a performance counter (col. 4 lines 1-13, "A counter tallies the clock ticks", wherein Applicant states on pg. 13 "that the performance counter 200 could be any type of programmable or re-settable counter", including counting a clock cycle).

As per claim 7, Reiffin teaches the method of claim 6 wherein the performance counter counts machine cycles in order to determine when to allocate the CPU resources (col. 4 lines 1-13, "A counter tallies the clock ticks", wherein each clock tick corresponds to a machine cycle).

As per claim 9, Reiffin teaches the method of claim 1 wherein the counter issues a first interrupt to the interrupt controller in order to instruct the interrupt controller to allocate the CPU resources (col. 4 lines 1-13, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU (central processing unit) to initiate an interrupt operation").

As per claim 10, Reiffin teaches the method of claim 9 wherein the interrupt controller issues a second interrupt to a CPU in order to instruct the CPU to allocate the CPU resources (col. 4 lines 1-13, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU (central processing unit) to initiate an interrupt operation").

As per claim 19, Reiffin teaches a computer-readable medium having computer-executable instructions stored for performing steps comprising:

a.   using a scheduler to control execution of at least one thread (Fig. 3);

b.   using at least one counter to notify the scheduler when to switch execution of said at least one thread (col. 4 lines 1-32, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU", "Control of the CPU is then passed to an interrupt service routine which contains or invokes the scheduler").

As per claim 20, Reiffin teaches the computer-readable medium of claim 19 further comprising instructions for issuing an interrupt from the counter when the counter reaches a predetermined number (col. 4 lines 1-13, "A counter tallies the clock ticks until the total reaches a predetermined number"), said predetermined number defining a maximum duration for execution of said at least one thread (col. 4 lines 1-13, "a predetermined number which determines the time duration of each timeslice"), said interrupt notifying the scheduler to switch execution to another thread (col. 4 lines 1-13, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU...to initiate an interrupt operation").

As per claim 21, Reiffin teaches the computer-readable medium of claim 20 wherein said at last one counter is a performance counter and counts CPU cycles (col. 4 lines 1-13, "A counter tallies the clock ticks", wherein the clock ticks correspond to CPU cycles).

As per claim 24, Reiffin teaches the computer-readable medium of claim 20 further comprising instructions for executing said at least one thread in a transparent manner so that at least one operating-system process is unaware of the execution of said at least one thread (col. 4

lines 41-59, "control of the CPU is rapidly switched back and forth between a local task and a network task so rapidly that the two tasks are said to be running concurrently", wherein the local task could be considered the operating system thread and can be treated simply as a permanent process, and does not need to be made aware of other background processes).

As per claim 25, Reiffin teaches the computer-readable medium of claim 24 further comprising instructions for executing all of said operating-system process and all of said at least one threads as a single real-time thread (col. 4 lines 41-59, "control of the CPU is rapidly switched back and forth between a local task and a network task so rapidly that the two tasks are said to be running concurrently", wherein since all the threads are running such that they appear to be running concurrently, they can all be considered a single real-time process).

As per claim 33, Reiffin teaches the computer-readable medium of claim 20 wherein said at least one thread and said another thread are of the same priority (col. 4 lines 33-40, wherein the scheduler chooses between a local task and a network task and the decision is made based purely on what the preempted task was. In that sense, the priorities of both tasks are equal since the type of task that was not running in ensured service when the other type of task is preempted).

Claim Rejections - 35 USC § 103

3. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negatived by the manner in which the invention was made.

4.   Claims 4-5, 8, 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over Reiffin in view of Patterson et al. (USPN 6,320,882) (hereinafter Patterson).

As per claims 4 and 5, Reiffin does not specifically teach the method of claims 2 and 3, respectively, wherein the CPU resources are allocated to at least one thread, and the CPU resources are allocated by determining a duration of time and a periodicity for execution of said at least one thread. However, Patterson does disclose such a concept (col. 2 lines 6-19, "each event requiring action to be taken after a predetermined time period, includes the steps of: generating values corresponding to each said event from a value of a global counter and the predetermined time period of the corresponding event", wherein the predetermined time period indicates the periodicity of execution for the event and the value including the counter indicates when the event expires, thus the duration of time).

It would have been obvious to one of ordinary skill in the art to combine Reiffin with Patterson since Patterson allows separate queuing of tasks based on their periodicity, thus providing more flexibility in the scheduling algorithm.

As per claim 8, Reiffin does not specifically teach the method of claim 6 wherein the performance counter counts executed computer instructions. Patterson does teach such an idea (col. 3 lines 41-43, "Global counter 222 is incremented based on the output of clock 204, and

may be, for example, incremented one 'tick' for every million cycles of clock 204" wherein this is merely one way in which the counter may be set up, and to align the clock such that a tick corresponds to the number of clock cycles necessary for a single instruction to issue is merely another embodiment of the same concept displayed by Patterson). The motivation for combining Patterson with Reiffin can be found above in reference to claims 4 and 5.

As per claim 22, Reiffin does not specifically teach the computer-readable medium of claim 20 wherein said at least one counter is a part of a CPU and counts executed instructions. Patterson does teach such an idea (col. 3 lines 41-43, "Global counter 222 is incremented based on the output of clock 204, and may be, for example, incremented one 'tick' for every million cycles of clock 204" wherein this is merely one way in which the counter may be set up, and to align the clock such that a tick corresponds to the number of clock cycles necessary for a single instruction to issue is merely another embodiment of the same concept displayed by Patterson). The motivation for combining Patterson with Reiffin can be found above in reference to claims 4 and 5.

5.   Claims 11-18 are rejected under 35 U.S.C. 103(a) as being unpatentable over Reiffin in view of Bonola (USPN 6,370,606).

As per claim 11, Reiffin does not specifically teach the method of claim 9 wherein the first interrupt is non-maskable. However, Bonola does teach the use of non-maskable interrupts (col. 10 lines 47-60, "the local APIC then checks the bus message to see whether it is a Non

Maskable Interrupt (NMI) (or an analogous signal)", wherein the interrupt controller is capable of handling non-maskable interrupts). It would have been obvious to one of ordinary skill in the art to combine Reiffin with Bonola since Reiffin does not specify the type of interrupt the interrupt controller handles from the counter. Therefore, by utilizing the method of Bonola a non-maskable interrupt would ensure that the interrupt is serviced and thus the real-time capabilities of the system are extended.

As per claim 12, Reiffin does not specifically teach the method of claim 10 wherein the first interrupt and second interrupt are non-maskable. However, Bonola does teach the use of non-maskable interrupts (col. 10 lines 47-60, "the local APIC then checks the bus message to see whether it is a Non Maskable Interrupt (NMI) (or an analogous signal)", wherein the interrupt controller is capable of handling non-maskable interrupts). The motivation for combining Bonola and Reiffin can be found above as discussed for claim 11.

As per claim 13, Reiffin teaches a method of scheduling resources on at least one microprocessor that includes a CPU and a device, the method comprising the steps of:

a.   using the device to determine when to allocate the resources (col. 4 lines 1-13, "A clock periodically emits timing pulses or ticks until the total reaches a predetermined number", wherein the counter tracks the ticks coming from the timing device and then tells the interrupt controller to allocate CPU resources);

b.   causing the device to issue an interrupt to the CPU when it is time to allocate the resources (col. 4 lines 1-13, "the counter...then signals the interrupt controller which in

turn activates the interrupt input of the CPU", wherein the timing device causes a chain reaction which interrupts the CPU when it is time to initiate an interrupt operation); and

c.   causing the CPU to allocate the resources in response to the interrupt (col. 4 lines 1-13, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU").

Reiffin does not specifically teach that the interrupt is non-maskable. Bonola does teach the use of a non-maskable interrupt (col. 10 lines 47-60, "the local APIC then checks the bus message to see whether it is a Non Maskable Interrupt (NMI) [or an analogous signal]"). The motivation for combining Reiffin with Bonola can be found above as discussed for claim 11.

As per claim 14, Reiffin teaches the method of claim 13 wherein the counter is a performance counter (col. 4 lines 1-13, "A counter tallies the clock ticks", wherein Applicant states on pg. 13 "that the performance counter 200 could be any type of programmable or re-settable counter", including counting a clock cycle).

As per claim 15, Reiffin teaches the method of claim 13 wherein the device is a timer (col. 4 line 1-13, "A clock periodically emits timing pulses or ticks at a high frequency", wherein the clock is analogous to a timer).

As per claim 16, Reiffin teaches a method of scheduling resources on at least one microprocessor that includes at least one performance counter, at least one programmable interrupt controller and at least one CPU, said method comprising the steps of:

a.   allowing the CPU to execute a first thread (col. 4 lines 1-13, wherein circumstances are described where an interrupt may be serviced, and thus the CPU would be executing a first task at the time of interruption);

b.   using the performance counter to determine when to allocate the resources to a second thread (col. 4 lines 1-13, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU", wherein when the interrupt controller activates the interrupt input of the CPU thereby relinquishing control from the first task to the second task);

c.   issuing a first interrupt from the performance counter to the programmable interrupt controller when it is time to allocate the resources to a second thread (col. 4 lines 1-13, "the counter...then signals the interrupt controller");

d.   instructing the programmable interrupt controller to issue a second interrupt to the CPU that instructs the CPU to switch execution from the first thread to the second thread (col. 4 lines 1-32, "the counter...then signals the interrupt controller which in turn activates the interrupt input of the CPU", "control of the CPU is taken away from the executing task which is thereby asynchronously preempted");

e.   instructing the CPU to stop execution of the first thread (col. 4 lines 14-32, "control of the CPU is taken away from the executing task which is thereby asynchronously preempted");

f.   causing the CPU to store first current state information regarding execution of the first thread ("Official Notice" is taken that it is well known and expected in the art to store state information of a preempted thread in a multithreaded environment, and since

Reiffin states that the executing task is preempted, it follows that it would have been obvious to store state information);

g.   causing the CPU to restore second current state information regarding execution of the second thread ("Official Notice" is taken that it is well known and expected in the art to restore state information of a previously suspended or blocked thread, and since Reiffin discusses multithreading, it follows that it would have been obvious to restore state information regarding the thread assuming control of the CPU); and

h.   allocating resources to the second thread (col. 4 lines 33-40, "The scheduler signals its task selection to the CPU which then normally executes either the scheduled local task or the scheduled network task", wherein once the next task to be scheduled is selected, CPU resources are allocated to that thread.

Reiffin does not specifically teach that the interrupt is non-maskable. Bonola does teach the use of a non-maskable interrupt (col. 10 lines 47-60, "the local APIC then checks the bus message to see whether it is a Non Maskable Interrupt (NMI) [or an analogous signal]"). The motivation for combining Reiffin with Bonola can be found above as discussed for claim 11.

As per claim 17, Bonola teaches the method of claim 16 wherein the programmable interrupt controller is an APIC (col. 10 lines 47-60, "the local APIC then checks the bus message to see whether it is a Non Maskable Interrupt (NMI) [or an analogous signal]").

As per claim 18, Bonola teaches the method of claim 17 wherein the microprocessor is selected from the group consisting of: a Pentium 4GB, a Pentium Pro 64GB, a Pentium MMX

4GB MMX, a Pentium II 4GB MMX, a Pentium II 4GB MMX KNI, a Celeron 4GB MMX, a Xeon PII 64GB MMX and a Xeon PIII 64GB MMX KNI (col. 11 lines 22-34, "four Pentium Pro-class processors A, B, C, and D...communicate amongst each other", wherein the above set of processors are applicable to the PC platform. Further, Bonola discusses the invention in relation to PC processors. Thus, the method of Bonola would be readily applicable to the other processors listed above).

6.   Claims 23 and 26-32 are rejected under 35 U.S.C. 103(a) as being unpatentable over Reiffin.

As per claim 23, Reiffin does not specifically teach the computer-readable medium of claim 20 further comprising instructions for executing said at least one thread at a highest IRQ level. However, "Official Notice" is taken that it would have been obvious to one of ordinary skill in the art to place the interrupting thread at a highest IRQ level since that would ensure that the interrupting task is serviced without being preempted by another task.

As per claim 26, Reiffin does not specifically teach the computer-readable medium of claim 20 wherein the duration for execution of said at least one thread is not equal to the duration for execution of said another thread. However, "Official Notice" is taken that varying execution durations are well known and expected in the art. Further, it would have been obvious to one of ordinary skill in the art that the duration a specific task must execute may vary and thus the

amount of time it requires the CPU would vary accordingly. As such, Reiffin discusses that the interrupting task should assume control of the CPU until the task has completed.

As per claim 27, Reiffin does not specifically teach the computer-readable medium of claim 19 further comprising instructions for allocation at least a portion of a CPU's resources to an operating-system process and using the remaining CPU resources for execution of said at least one thread. However, "Official Notice" is taken that the use of an operating system thread is well known and expected in the art of multithreading. In such a case, when an operating system thread exists, all other threads run at a lower priority and compete for the remaining CPU resources. For an example of such an idea, see the attached reference Kosche et al. (USPN 5,937,187).

As per claim 28, Reiffin does not specifically teach the computer-readable medium of claim 27 further comprising instructions for releasing the CPU resources back to the operating-system process when said at least one thread finishes execution. However, "Official Notice" is taken that it is well known that operating systems handle the CPU resources and distribute them among tasks. As discussed for claim 27, operating system threads are well known in the art. Thus, the managing of CPU resources, such as reclaiming CPU resources when a thread completes, is an inherent function of an operating system thread.

As per claim 29, Reiffin teaches the computer-readable medium of claim 27 further comprising instructions for releasing the CPU resources to another thread when said at least one

thread finished execution (col. 4 lines 33-40, "At the start of the next timeslice the counter is reset and the above-described cycle of operation is iterated over and over again", wherein the cycle of servicing tasks through the use of interrupts is taught, and upon completion of a timeslice, the yielding thread returns the CPU resources to the system).

As per claim 30, it is rejected for similar reasons as stated for claim 27.

As per claim 31, Reiffin does not specifically teach the computer-readable medium of claim 20 further comprising instructions for switching the interrupt from non-maskable to maskable. However, Applicant discloses Prior Art that states that maskable interrupts can be enabled by executing the STI instruction. Thus, the Prior Art already suggests the act of changing an interrupt from non-maskable to maskable.

As per claim 32, it is rejected for similar reasons as stated for claim 31.

Claim Objections

7.   Claim 32 objected to under 37 CFR 1.75(c), as being of improper dependent form for failing to further limit the subject matter of a previous claim. Applicant is required to cancel the claim(s), or amend the claim(s) to place the claim(s) in proper dependent form, or rewrite the claim(s) in independent form. Claim 32 is a duplicate of claim 31.
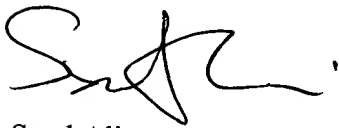
Conclusion

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Syed J Ali whose telephone number is (703) 305-8106. The examiner can normally be reached on Mon-Fri 8-5:30, 1st Friday off.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John A Follansbee can be reached on (703) 305-8498. The fax phone numbers for the organization where this application or proceeding is assigned are (703) 746-7239 for regular communications and (703) 746-7238 for After Final communications.

Any inquiry of a general nature or relating to the status of this application or proceeding should be directed to the receptionist whose telephone number is (703) 305-3900.

Syed Ali
March 11, 2003

MAJID BANANKHAH
PRIMARY EXAMINER